Asymptotically optimal prior-free asset market mechanisms[*]

Simon Loertscher[†] Leslie M. Marx[‡]

February 8, 2021

Abstract

We develop a prior-free mechanism for an asset market that is dominant-strategy incentive compatible, ex post individually rational, constrained efficient, and asymptotically optimal—the profit of the market maker using this mechanism approaches the profit that it would make using the optimal mechanism for the setting in which the designer knows the distribution of the agents' types. The mechanism is direct. It first identifies agents with values equal to the Walrasian price. As the number of agents goes to infinity, these will almost surely not trade under the Bayesian optimal mechanism. The second step can be described algorithmically as consisting of ascending and descending clock auctions that start from the Walrasian price and estimate virtual types along the way. The mechanism thus permits partial clock auction implementation. It can accommodate heterogeneity among groups of traders and discrimination among these, provided heterogeneity is not too accentuated.

Keywords: market mechanism, endogenous trading position, detail free, mechanism design with estimation

JEL Classification: C72, D44, L13

---

[*]We are grateful for valuable discussion with Bing Liu and Ellen Muir, and we thank seminar participants at the University of Melbourne. Financial support from the Samuel and June Hordern Endowment and the Australian Research Council under a Discovery Project Grant DP200103574 is gratefully acknowledged.

[†]Department of Economics, Level 4, FBE Building, 111 Barry Street, University of Melbourne, Victoria 3010, Australia. Email: simonl@unimelb.edu.au.

[‡]Fuqua School of Business, Duke University, 100 Fuqua Drive, Durham, NC 27708, USA: Email: marx@duke.edu.

1 Introduction

Exchanges in which agents' trading positions—buy, sell, or hold—are determined endogenously abound in the real world. If anything, they have become even more prevalent in the digital age as companies like Uber, Lyft, and AirBnB allow agents who used to be active on only one side of the market as buyers or sellers to choose whether they consume or provide services. Moreover, vertical integration almost inevitably makes the integrated firm's trading position endogenous since, for some realizations of its own values and costs and those of outsiders, it may optimally buy additional inputs, sell to the outside market, or decide not to trade. While the mechanism design literature has long paid scant attention to such markets, which we refer to as *asset markets*, there has been a recent upsurge of interest.[1] An important issue, of particular salience in novel industries and environments such as those mentioned concerns the mechanism for an intermediary that faces a large number of traders without knowing the distributions of their types.

In this paper, we address the question: is there a mechanism that satisfies dominant-strategy incentive compatibility, respects the agents' individual rationality constraints ex post, and converges, as the number of agents who independently draw their types, endowments, and demands grows large, to the mechanism that would maximize the designer's profit if the designer knew the distribution from which the agents draw their types? For an environment with commonly known endowments and maximum demands and independently distributed, constant marginal values, we show that the answer is affirmative if the (unknown) distribution is regular in the sense of exhibiting both increasing virtual cost and values. Moreover, we show that there is such an *asymptotically optimal* prior-free mechanism that, in addition, satisfies *constrained efficiency* insofar as the agents who sell their endowments under this mechanism are the ones with the lowest values and the agents who buy under this mechanism are the ones with the highest values. Constrained efficiency is a desirable property because it eliminates post-allocation gains from trade on each side of the market and thereby reduces scope for resale and, related, bid shading and public fall-out due to discrimination that arises when, say, an agent who submitted a lower bid obtains a unit while another with a higher bid is not served; see Loertscher and Mezzetti

[1]See, for example, Lu and Robert (2001), Chen and Li (2018), Loertscher and Wasser (2019), Johnson (2019), Loertscher and Marx (2020b), Delacrétaz et al. (2020) and Li and Dworczak (2020), who study models along these lines. The labels may differ from one paper to another.

(forth.). In light of the Wilson critique (Wilson, 1987), being prior-free is a desirable robustness property, while asymptotic optimality ensures the designer against ex post regret—the designer would not have made more profit, ex post, had it known at the outset the type distribution that it learned by running the mechanism (Loertscher and Marx, 2020a).

A unique feature of asset markets and a major challenge to designing suitable mechanisms for asset market environments is that they are subject to what Lewis and Sappington (1989) dubbed "countervailing incentives": agents who are likely to trade as sellers have incentives to exaggerate their costs, while agents with high values are likely to trade as buyers and therefore have incentives to report lower values than they have. This implies that in an asset market, an agent's information rent is typically minimized for a type in the interior of the type space. This contrasts sharply with one-sided allocation problems such as auctions or procurement auctions and the two-sided allocation problems addressed by double auctions, where it is a priori known which agents are buyers and sellers and therefore a priori known which type of an agent is the worst-off—the lowest possible type for a buyer and the highest possible type for a seller.

From the designer's perspective, knowledge of which types are worst-off identifies the types for which the individual rationality constraint binds, which, in turn, permits clock auction implementation of the optimal mechanism in environments with one-sided and two-sided private information. Because of the countervailing incentives and the endogeneity of trading positions inherent in asset markets, these markets afford no simple, a priori known point at which one can, say, start a clock auction. The mechanism we develop is, therefore, direct. It first identifies agents with values equal to the Walrasian price defined with respect to the reported types. As the number of agents goes to infinity, these agents will almost surely not trade under the mechanism that is optimal for the intermediary in the same sense and under the same conditions as Myerson's optimal auction is optimal for the auctioneer. Its second step can be described algorithmically as consisting of ascending and descending clock auctions that start from the Walrasian price and estimate virtual types along the way using the methodology developed in Loertscher and Marx (2020a).

Although the trades that are executed are efficient in that they involve sales from the lowest-type agents to the highest-type agents, the mechanism that maximizes the market maker's expected profit does not, in general, involve the efficient number of trades. Just as a monopolist sells less than the efficient quantity, our market maker generally restricts the quantity of trades below the efficient level. In contrast to an optimal sales (optimal procurement) auction, the distortion away from the efficient allocation does not involve traders with low (high) values but rather agents of intermediate types equal to or near the Walrasian price. These are the types for whom the individual rationality constraint is binding in an optimal mechanism.

As we show, our mechanism permits partial clock implementation, which is desirable to the extent, as argued, for example, by Baranov et al. (2017), that even partial privacy preservation is valuable when there is a tension between privacy preservation and price discovery. For example, one can start an ascending clock. The clock increases as long as indicated net demand exceeds supply and agents that exit either do not trade or trade as a seller at a price that does not depend on its bid (other than its having exited). As soon as net demand is less than supply, the mechanism begins estimating virtual values and virtual costs, beginning in both directions from the price at which net demand equals supply. The clock auction continues to reduce the quantity demanded until estimated virtual values and estimated virtual costs are the same. This ascending clock auction preserves the privacy of agents who trade as buyers and endows agents with obviously dominant strategies in the subgame that ensues after the Walrasian price is reached.[2] Conversely, one can implement the mechanism using a descending clock auction and preserve the privacy of the agents who trade as sellers.

The main analysis assumes that all agents draw their types from the same distribution and have the same endowments and maximum demands, where for simplicity we set the ratio of maximum demand over endowment equal to two. In an extension, we show that the assumptions on maximum demands and endowments can be relaxed straightforwardly. We also provide conditions under which one can allow for heterogeneous groups of traders, with traders in each group drawing their types from the same regular distribution, while the distributions are allowed to vary across groups.

To the best of our knowledge, the mechanism design literature on asset markets is confined to the papers mentioned in footnote 1, none of which analyzes prior-free asymptotic optimality. Outside the domain of asset markets and beginning with Baliga and Vohra (2003) and Segal (2003), there is a sizeable literature on prior-free asymptotic optimal mechanisms. To maintain incentive compatibility when agents

[2]See Li (2017) for the definition of obviously dominant strategies.

compete against each other for the opportunity to trade, most of this literature takes either a random-sampling approach, whereby a small subset of agents is sampled and prevented from trading, or splits the market into separate submarkets and uses reports and estimates from one submarket to determine the allocation and payments in the other. Neither of these mechanisms is constrained efficient. An exception is Loertscher and Marx (2020a), which develops a clock auction that is prior free and asymptotically optimal if the type distributions satisfy regularity; we refer the reader to that paper for additional references.

The remainder of this paper is organized as follows. Section 2 introduces the setup. In Section 3, we introduce the theoretical benchmark that is used to evaluate the performance of our prior-free mechanism, and we define asymptotic optimality. In Section 4, we derive asymptotically optimal, prior-free mechanisms. Section 5 contains extensions, and Section 6 concludes the paper.

2 Setup

We assume that there is a set $\mathcal{N}$ of $n \geq 3$ agents, where n is odd. Each agent $i \in \mathcal{N}$ has an endowment of one unit of a homogeneous good and a constant marginal value θ_i for up to two units of the good, which we also refer to as the agents' capacities. The willingness to pay for additional units beyond the capacity is zero.

The agents' endowments and capacities are commonly known, including by the market maker (or designer). Agents' types are their own private information and are drawn independently from a continuously differentiable distribution F with compact interval support, which we normalize without loss of generality to $[0, 1]$, and density f that is positive on $[0, 1]$. To simplify notation, we assume that there are no ties among the agents' types.[3] We assume further that F has increasing virtual type functions

$$\Phi(\theta) \equiv \theta - \frac{1 - F(\theta)}{f(\theta)} \quad \text{and} \quad \Gamma(\theta) \equiv \theta + \frac{F(\theta)}{f(\theta)},$$

where we refer to Φ as the virtual value function and to Γ as the virtual cost function.

We assume that the market maker does not know the specifics of F, but does know that F has a continuous, positive density on $[0, 1]$ and increasing virtual type

[3] One can accommodate ties by introducing an ordering over agents as in Loertscher and Marx (2020b) by using that ordering to define the tie breaking rule.

functions. We assume that the agents also do not know F, so that the market maker cannot simply ask them to report it.

We say a mechanism is *prior-free* if it is defined without reference to the type distribution. A mechanism is *asymptotically optimal* if it is dominant-strategy incentive compatible and ex post individually rational and has the property that the ratio of the market maker's expected profit in the mechanism relative to its expected profit in the Bayesian optimal mechanism converges to 1 in probability as n goes to infinity (with n odd).

This setup provides a useful starting point for the analysis because the median-type agent does not trade under efficiency and so also does not trade under profit maximization. Profit maximization comes down to the selection of which agents should sell their endowments and which agents should buy those units. While the underlying economics is most transparent in this simple setup, the assumptions that endowments and maximal demands are the same across all agents (and equal to one and two, respectively) and that the number of agents is odd are inessential, as we show in Section 5.3. As discussed at the end of Section 4.1, the assumption that n is odd simplifies the analysis without being essential.

3 Optimal mechanisms

In this section, we first provide the theoretical benchmark against which the performance of the prior-free mechanism will be measured.

3.1 Bayesian optimal mechanism

The market maker's objective is to maximize its expected profit,[4] which, as we shall see, may imply a smaller number of trades than under efficiency. For a given allocation rule $\mathbf{s} : \boldsymbol{\theta} \to \{0,1,2\}^n$ with $\sum_{i \in \mathcal{N}} s_i(\boldsymbol{\theta}) = n$, we define the profit to the designer in terms of the dominant-strategy implementation: each buyer pays and each seller receives its threshold type for each unit that it trades, where a buyer's threshold type is the lowest type that the buyer could report and still trade, and a seller's threshold type is the highest type that the seller could report and still trade.

[4] It is straightforward to allow the market maker's objective to be the maximization of the expected value of a weighted average of profit and social surplus, as described in Section 5.1.

We define a *Bayesian optimal* mechanism to be a mechanism that maximizes the market maker's expected profit, subject to incentive compatibility and individual rationality, given knowledge of the type distribution.

Because agents can be either buyers or sellers, the Bayesian optimal mechanism is defined in terms of ironed virtual types that combine the virtual cost and virtual value functions. Specifically, let $\overline{x}$ be such that $\Gamma(\overline{x}) = \Phi(\overline{\theta}) = \overline{\theta}$, and for any $x \in [\underline{\theta}, \overline{x}]$, let $y(x)$ be such that $\Phi(y(x)) = \Gamma(x)$. Then for any ironing region $[x, y(x)]$, where $x \in [\underline{\theta}, \overline{x}]$, define the ironed virtual type function as

$$V(\theta \mid x) \equiv \begin{cases} \Gamma(\theta) & \text{if } \theta < x, \\ \Gamma(x) & \text{if } \theta \in [x, y(x)], \\ \Phi(\theta) & \text{if } y(x) < \theta, \end{cases}$$

which is illustrated in Figure 1. It follows from our regularity assumptions that $V(\theta \mid x)$ is nondecreasing in θ.

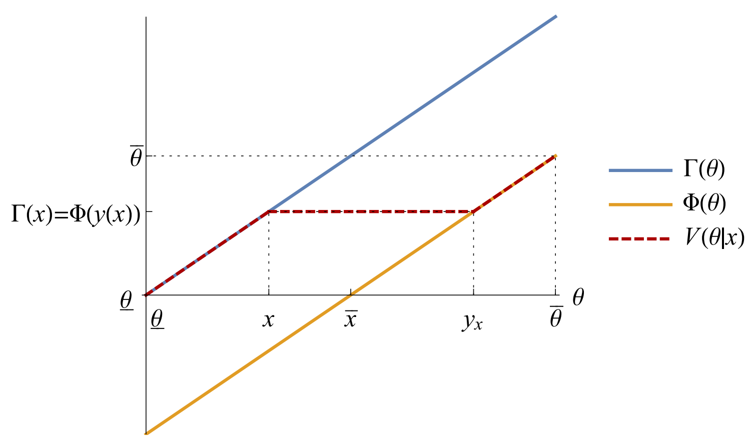

Figure 1: Illustration of the ironed virtual type function.

Lu and Robert (2001) show that for finite n, the Bayesian optimal allocation rule and ironing region are such that the allocation rule maximizes the expected value of the sum of ironed virtual types, weighted by quantities, and such that the interim expected trade for agents with types in the ironing region is zero. Adapting Lu and Robert's result to our setup, and letting $\theta_{(j)}$ denote the j-th highest type among $\theta_1, \ldots, \theta_n$ and $\theta_{[j]}$ denote the j-th lowest type, with $\theta_{(0)} \equiv 1$ and $\theta_{[0]} \equiv 0$, we have the following result:

Proposition 1 *The Bayesian optimal outcome is defined by* $x^* \in [\underline{\theta}, \overline{x}]$ *and* $q : [\underline{\theta}, \overline{\theta}]^n \to \{0, \ldots, (n-1)/2\}$ *such that* q *solves*

$$\max_{q:[\underline{\theta},\overline{\theta}]^n \to \{0,\ldots,(n-1)/2\}} \mathbb{E}_{\boldsymbol{\theta}} \left[\sum_{i=1}^{q(\boldsymbol{\theta})} (V(\theta_{(i)} \mid x^*) - V(\theta_{[i]} \mid x^*)) \right],$$

and the interim expected trade for an agent of type $\theta \in [x^*, y(x^*)]$ *is zero, i.e., for all* $\theta_i \in [x^*, y(x^*)]$,

$$\Pr_{\boldsymbol{\theta}_{-i}}(\theta_i \geq \theta_{(q(\boldsymbol{\theta}))}) - \Pr_{\boldsymbol{\theta}_{-i}}(\theta_i \leq \theta_{[q(\boldsymbol{\theta})]}) = 0.$$

3.2 Limit properties of the Bayesian optimal mechanism

As we now show, the extension of Lu and Robert (2001, Theorem 3) to the limit case in our setup can be defined simply in terms of seller and buyer prices p^S and p^B that satisfy a market clearing condition,

$$1 - F(p^B) = F(p^S), \tag{1}$$

and an optimality condition,

$$\Phi(p^B) = \Gamma(p^S), \tag{2}$$

as follows:

Proposition 2 *In the limit as* n *goes to infinity, with probability one, the Bayesian optimal outcome has each agent with type less than* p^S *sell one unit and each agent with type greater than* p^B *buy one unit.*

Proof. See Appendix B.

Proposition 2 implies that the market maker's limit payoff per agent is

$$(p^B - p^S)F(p^S),$$

where p^S and p^B are defined by (1) and (2). For example, if types are drawn from the uniform distribution, then $\Phi(\theta) = 2\theta - 1$ and $\Gamma(\theta) = 2\theta$, so we have $p^B = 3/4$, $p^S = 1/4$ and $F(p^S) = 1/4$. In the limit, the designer's per-agent expected profit is 1/8.

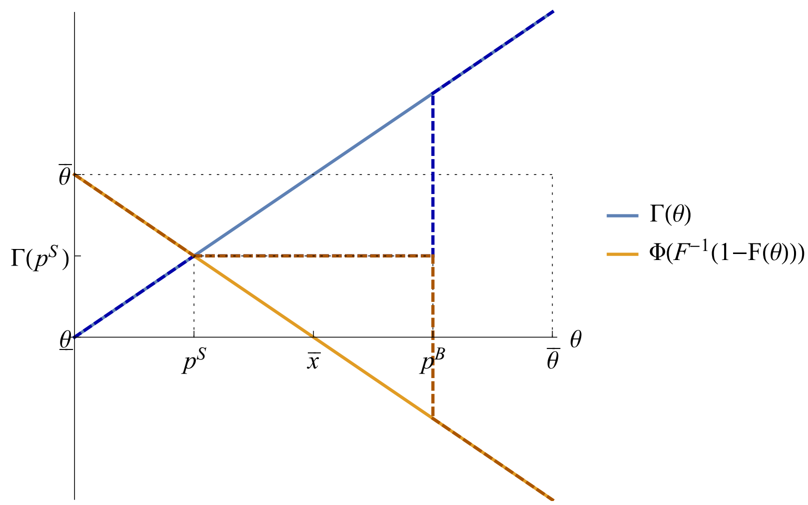

Figure 2: Illustration that in the limiting case as n goes to infinity, the virtual cost function Γ and virtual value function Φ "cross" at the boundary of the ironing region, $[p^S, p^B]$.

As described in the proof of Proposition 2 and shown in Figure 2, in the limiting case as n goes to infinity, the virtual cost and virtual value functions "cross" at the boundary of the ironing region, $[p^S, p^B]$. Thus, in the limit, it is a probability zero event that there are trades by agents with types in the ironing region. As a result, in developing an asymptotically optimal prior-free mechanism, we can focus on estimating the virtual values and virtual costs, without having to also estimate the ironing region. This turns out to be critical because the virtual value and virtual cost for the marginal traders can be estimated based on the types of nontrading agents; however, the estimation of the ironing region requires information about the types of trading agents.

As Proposition 2 shows, knowledge of the distribution F and associated virtual type functions is essential for the design of the Bayesian optimal mechanism. The question then remains how a market marker can design a mechanism in the absence of such knowledge, which is what we address in the next section.

4 Asymptotically optimal prior-free mechanisms

We begin by defining a prior-free, incentive compatible, individually rational asset market mechanism with estimated virtual types, and we then define the prior-free estimators for the virtual types to be used in that mechanism. We conclude the section by showing that the prior-free asset market mechanism with estimated virtual types can be implemented using a partial clock auction that preserves the privacy of trading agents on one side of the market.

4.1 Prior-free mechanism with estimated virtual types

Although the prior-free asset market mechanism with estimated virtual types that we define is a direct mechanism that requires agents to report their types only once, at the outset, one can usefully think of it as consisting of two phases.

In the first phase, given reports $\theta_1, \ldots, \theta_n$, the mechanism determines the type of the agent that, under efficiency, does not trade: $\hat{\theta} \equiv \theta_{\left(\frac{n+1}{2}\right)}$. We refer to $\hat{\theta}$ as the *cutoff type* and the agent with type $\hat{\theta}$ as the *cutoff agent*. Note that $\hat{\theta}$ is the (unique) market clearing price given the realization $\boldsymbol{\theta}$. After the cutoff type is identified, the mechanism has a second phase, the "auction phase," in which agents with initial reports above the cutoff type participate as buyers in an ascending clock auction starting at the cutoff type, while agents with reports below the cutoff type participate as sellers in a descending clock auction. In this auction phase, for $t \in \{0, 1, \ldots\}$,[5] we have a state $\boldsymbol{\omega}_t = (\mathcal{B}_t, \mathcal{S}_t, p_t^b, p_t^s)$, where the components of the state are: the set of active buyers $\mathcal{B}_t \subseteq \mathcal{N}$, the set of active sellers $\mathcal{S}_t \subseteq \mathcal{N}$, the buyer price $p_t^b \in [0, 1]$, and the seller price $p_t^s \in [0, 1]$. Letting Ω be the set of all such states, the auction phase relies on a virtual value estimator $\phi : \Omega \to \mathbb{R}$ that is increasing in the buyer price and a virtual cost estimator $\gamma : \Omega \to \mathbb{R}$ that is increasing in the seller price.

We initialize the auction phase by setting $\mathcal{S}_0 \equiv \{i \in \mathcal{N} \mid \theta_i < \hat{\theta}\}$, $\mathcal{B}_0 \equiv \{i \in \mathcal{N} \mid \theta_i > \hat{\theta}\}$, and $p_0^b = p_0^s \equiv \hat{\theta}$. The cutoff agent is not included in either of the sets $\mathcal{S}_0$ or $\mathcal{B}_0$—as a result, the cutoff agent never trades. Also, $|\mathcal{B}_0| = |\mathcal{S}_0| = \frac{n-1}{2} \geq 1$, so, given that $n \geq 3$, the auction starts with at least one active buyer and seller.

For $t \in \{0, 1, \ldots\}$, the auction phase applies the following iterative process:

- Step 1: If either $\phi(\boldsymbol{\omega}_t) \geq \gamma(\boldsymbol{\omega}_t)$ or $|\mathcal{B}_t| = 1$, then the auction phase ends; otherwise, proceed to step 2.

- Step 2: Let
$$\boldsymbol{\omega}_{t+1} = \left(\mathcal{B}_t \backslash \{j^B\},\ \mathcal{S}_t \backslash \{j^S\},\ \theta_{j^B},\ \theta_{j^S}\right),$$

[5]The mechanism that we define has no more that $\frac{n-3}{2}$ iterations, so t never exceeds this value.

where j^B and j^S are defined by $\theta_{j^B} = \min\{\theta_i\}_{i \in \mathcal{B}_t}$ and $\theta_{j^S} = \max\{\theta_i\}_{i \in \mathcal{S}_t}$, and proceed to step 3.

- Step 3: Increment t by 1 and return to step 1.

Because step 2 reduces the number of active buyers by 1 and the number of active sellers by 1, the auction phase ends after a finite number of iterations. Let $\hat{\boldsymbol{\omega}} = (\hat{\mathcal{B}}, \hat{\mathcal{S}}, \hat{p}^b, \hat{p}^s)$ denote the state when the auction phase ends. By construction, $|\hat{\mathcal{B}}| = |\hat{\mathcal{S}}|$, $\hat{p}^b \geq \hat{\theta}$, and $\hat{p}^s \leq \hat{\theta}$. The asset market mechanism then specifies that $|\hat{\mathcal{B}}|$ units are traded, with each agent $i \in \hat{\mathcal{B}}$ receiving one unit and paying $\hat{p}^b$, and each agent $i \in \hat{\mathcal{S}}$ providing its endowment and being paid $\hat{p}^s$.

Proposition 3 *The prior-free asset market mechanism with virtual type estimators ϕ and γ is dominant-strategy incentive compatible, individually rational, and deficit free if for any state $\boldsymbol{\omega} = (\mathcal{B}, \mathcal{S}, p^B, p^S)$, $\phi(\boldsymbol{\omega})$ and $\gamma(\boldsymbol{\omega})$ depend only on p^B, p^S, and the types of agents in $\mathcal{N} \backslash (\mathcal{B} \cup \mathcal{S})$.*

Proof of Proposition 3. If $\phi(\boldsymbol{\omega})$ and $\gamma(\boldsymbol{\omega})$ depend only on p^B and p^S (which are functions of the types of agents in $\mathcal{N} \backslash (\mathcal{B} \cup \mathcal{S})$) and on the types of agents in $\mathcal{N} \backslash (\mathcal{B} \cup \mathcal{S})$, then the stopping rule in the auction phase depends only on the types of agents that do not trade and so the auction phase is a two-sided clock auction as defined by Loertscher and Marx (2020a).[6] By the nature of clock auctions, it follows that in the auction phase every agent has a dominant strategy to bid truthfully. Hence, we are left to rule out profitable deviations from truthful reporting in the first phase, which determines whether an agent is the cutoff agent (and so does not trade) or participates as a buyer or a seller.

Consider first an agent who under truthful reporting is the cutoff agent. If the agent deviates from truthful reporting but remains the cutoff agent, its payoff remains zero and so the deviation is not profitable. If the agent reports a higher type so that it becomes a buyer, the cutoff type increases (above the deviating agent's type), so that if the agent trades in the clock auction, its payoff is negative. The best case for the agent is that it does not trade in the clock auction, in which case its payoff is zero. Thus, a deviation that causes the agent to become a buyer in the clock auction

[6] The two-sided clock auction defined in Loertscher and Marx (2020a) is an adaptation of the one-sided clock auction of Milgrom and Segal (2020).

is not profitable. A symmetric argument shows that a deviation that causes the agent to become a seller in the clock auction is also not profitable.

We are left to consider agents who under truthful reporting are buyers or sellers in the clock auction. The smallest payoff from truthful reporting in the clock auction is zero, so the deviation to become the cutoff agent is not profitable. Any deviation that leaves an agent on the same side of the cutoff agent does affect the cutoff type, and so has no effect on the starting prices of the clock auctions. Thus, no such deviation is profitable.

The only deviations we are left to rule out as profitable are those such that an agent who under truthful reporting is a buyer (seller) becomes a seller (buyer). But no such deviation can be profitable because the best that could happen to the deviating agent is that it does not trade, in which case its payoff is no more than the minimum payoff from truthful reporting in the clock auction. If the agent does trade, it either sells at a price below its true type or buys at a price above its true type, in which case it makes a loss.

Thus, the mechanism is dominant-strategy incentive compatible. That it is individually rational follows from the fact by bidding truthfully that every agent obtains a net payoff of at least zero. Deficit freeness follows from the facts that the lowest price paid by any buyer on any unit traded is the cutoff type, while the highest price received by any seller on any unit is the cutoff type, and that the quantity traded is balanced. ■

Proposition 3 puts us in a position to construct an asymptotically optimal prior-free market mechanism, provided that we have mappings ϕ and γ that give consistent estimates of the marginal buyer's virtual value and the marginal seller's virtual cost, respectively, without interfering with the incentive compatibility of the mechanism. The critical issue, then, is to construct consistent, prior-free estimates of the hazard rates $(1-F(\theta))/f(\theta)$ and $F(\theta)/f(\theta)$ using only data from agents that do not trade. As we show below, the spacings-based estimator of Loertscher and Marx (2020a) provides such an estimator.

While the analysis above assumes that n is odd, one can easily accommodate an even number of agents n with $n \geq 4$. To do so, define the two "middle" types, $\hat{\theta}_1 \equiv \theta_{(n/2+1)}$ and $\hat{\theta}_2 \equiv \theta_{(n/2)}$, where $\hat{\theta}_1 \leq \hat{\theta}_2$, and initialize the auction phase with $p_0^S \equiv \hat{\theta}_1$, $p_0^B \equiv \hat{\theta}_2$, $\mathcal{S}_0 \equiv \{i \in \mathcal{N} \mid \theta_i < \hat{\theta}_1\}$, and $\mathcal{B}_0 \equiv \{i \in \mathcal{N} \mid \theta_i > \hat{\theta}_2\}$.

4.2 Asymptotic optimality

We now specify virtual type estimators ϕ_σ and γ_σ that satisfy the conditions of Proposition 3 and that have properties sufficient to guarantee asymptotic optimality.[7] Guided by Loertscher and Marx (2020a), for any state $\boldsymbol{\omega} = (\mathcal{B}, \mathcal{S}, p^B, p^S)$ with the number of active agents on each side denoted by $a \equiv |\mathcal{B}| = |\mathcal{S}| \in \{1, \ldots, \frac{n-1}{2}\}$, we use

$$\phi_\sigma(\boldsymbol{\omega}) \equiv p^B - (a+1)\,\sigma^B(\boldsymbol{\omega}) \text{ and } \gamma_\sigma(\boldsymbol{\omega}) \equiv p^S + (a+1)\,\sigma^S(\boldsymbol{\omega}), \tag{3}$$

where $\sigma^B(\boldsymbol{\omega})$ and $\sigma^S(\boldsymbol{\omega})$ are spacing estimators given by

$$\sigma^B(\boldsymbol{\omega}) \equiv \begin{cases} \frac{\theta_{(a+1)} - \theta_{(a+1+d(\boldsymbol{\omega}))}}{d(\boldsymbol{\omega})} & \text{if } d(\boldsymbol{\omega}) > 0, \\ \frac{1}{n+1} & \text{otherwise,} \end{cases} \quad \sigma^S(\boldsymbol{\omega}) \equiv \begin{cases} \frac{\theta_{[a+1+d(\boldsymbol{\omega})]} - \theta_{[a+1]}}{d(\boldsymbol{\omega})} & \text{if } d(\boldsymbol{\omega}) > 0, \\ \frac{1}{n+1} & \text{otherwise,} \end{cases} \tag{4}$$

and where $d(\boldsymbol{\omega})$ is the number of spacings to be included in the average, defined by

$$d(\boldsymbol{\omega}) \equiv \min\left\{\lceil n^b \rceil, n - 2a - 1\right\} \tag{5}$$

for some $b \in (0,1)$, where $\lceil n^b \rceil$ is the smallest integer greater than or equal to n^b.[8]

Our prior-free asset market mechanism can therefore be summarized as $\langle \phi_\sigma, \gamma_\sigma, b \rangle$, defined by (3)–(5), where $b \in (0,1)$. Unpacking this, given a state $\boldsymbol{\omega}$ such that a active buyers and a active sellers remain, $\sigma^B(\boldsymbol{\omega})$ is the average spacing between the $d(\boldsymbol{\omega})$ highest types for nontrading agents, i.e., the average spacing between the $d(\boldsymbol{\omega})$ highest types among

$$\theta_{[a+1]}, \ldots, \theta_{\left[\frac{n+1}{2}\right]} = \hat{\theta} = \theta_{\left(\frac{n+1}{2}\right)}, \ldots, \theta_{(a+1)},$$

and analogously for the seller-side spacing $\sigma^S(\boldsymbol{\omega})$. The number of spacings to include in the averages obviously cannot exceed the number of observed spacings, which is $n - 2a - 1$.[9] In addition, the number of spacings to include is capped at n^b with $b \in (0,1)$ so that as n grows large, the number of spacings that can be included in

[7] If $n = 3$, then the auction phase ends at $t = 0$ with 1 active buyer and 1 active seller, who trade at a price of $\hat{\theta}$, so in that case, the virtual type estimators do not play a role.

[8] The choice of $b = 4/5$ minimizes the mean squared error of nearest neighbor estimators such as σ^B and σ^S (Silverman, 1986, Chapters 3 and 5.2.2). For related discussion, see Loertscher and Marx (2020a).

[9] During the auction phase at $t = 0$, the sets of active buyers and sellers have $\frac{n-1}{2}$ members, so $d(\boldsymbol{\omega_0}) = 0$. Thus, if $n \geq 5$, one has at most $\frac{n-3}{2}$ trades, even though having $\frac{n-1}{2}$ trades might increase the designer's expected payoff.

the average grows large in absolute terms, but shrinks as a proportion of the total, i.e., as n goes to infinity, n^b goes to infinity, but n^b/n goes to zero.[10] In addition, the share of spacings included in the average does not shrink too quickly in the sense that $\lim_{n\to\infty} \frac{\ln n}{n^b} = 0$. As a result, given our continuity assumptions on the distribution of types, the spacing estimators $\sigma^B(\boldsymbol{\omega})$ and $\sigma^S(\boldsymbol{\omega})$ are consistent estimators for the spacing between types close to $\theta_{(a)}$ and $\theta_{[a]}$, respectively, and we have the following result:

Proposition 4 *The prior-free asset market mechanism* $\langle \phi_\sigma, \gamma_\sigma, b\rangle$ *with* $b \in (0,1)$ *is asymptotically optimal.*

Proof. See Appendix B.

While Proposition 4 shows that prior-free asset market mechanism $\langle \phi_\sigma, \gamma_\sigma, b\rangle$ with $b \in (0,1)$ is asymptotically optimal, it leaves open the question as to what value b should take. The tradeoff is that with a smaller value of b, the hazard-rate estimator uses less information, but that information is more tightly clustered around the point where the hazard rate is being estimated; a larger value of b uses more information, but risks using information that is less relevant for the estimation of the hazard rate at the point under consideration. As discussed in Loertscher and Marx (2020a), using $b = 4/5$ gives the estimator that minimizes the mean square error of the estimator.[11]

While the prior-free asset market mechanism $\langle \phi_\sigma, \gamma_\sigma, b\rangle$ is optimal in the limit as the number of agents grows large, it can, of course, be used in settings with a finite, and even small, number of agents. For a given b, if n is sufficiently small, then all non-trading agents types are used in estimating the hazard rates. In that case, the number of agents that trade is independent of the type realizations.

Proposition 5 *For any* $b \in (0,1)$, *if* $2\left\lceil \frac{n+1}{4} \right\rceil \leq n^b$, *then the prior-free asset market mechanism* $\langle \phi_\sigma, \gamma_\sigma, b\rangle$ *has* $\frac{n-1}{2} - \left\lceil \frac{n+1}{4} \right\rceil$ *trades, regardless of the type realizations.*

Proof. See Appendix 5.

[10] Because all agents, including those that ultimately become buyers and those that ultimately become sellers, draw their types from the same distribution, we can use the spacings on both the buyer and seller sides, i.e., all available $n - 2a - 1$ spacings. That said, for n sufficiently large, the n^b cap will bind sufficiently that we will only be using buyer-side spacings for the buyer-side spacing estimator and seller-side spacings for the seller-side spacing estimator.

[11] For foundations for this, see Silverman (1986, Chapters 3 and 5.2.2).

For example, Proposition 5 implies that if $\lceil \frac{n+1}{4} \rceil \leq n^b$ and $n+1$ is an integer multiple of 4, then the prior-free asset market mechanism $\langle \phi_\sigma, \gamma_\sigma, b \rangle$ has $\frac{n+1}{4} - 1$ trades. Roughly speaking, in the small-market implementation of the prior-free asset market mechanism, the agents with types in the top 12.5% of types purchase from the agents in the bottom 12.5%, with the middle 75% of agents not trading.

4.3 Partial clock auction implementation

It is possible to implement the mechanism using a clock auction that applies to only one side of the market, thereby preserving the privacy of all trading agents on one side of the market. An ascending clock preserves the privacy of all traders who end up as buyers (i.e., whose types are larger than the market clearing price), and a descending clock auction preserves the privacy of all analogously defined sellers. In this way, the mechanism strikes a balance between eliciting the information required to operate and preserving the privacy of agents where possible. This is in a similar vein to the one-sided auction of Baranov et al. (2017), which induces losing bidders to reveal cost information about quantities that are no longer profitable, but generally does not require winning bidders to reveal their cost for the awarded quantities. As they remark, "the format strikes a balance between price discovery and privacy preservation" (Baranov et al., 2017, p. 5).

In contrast to a one-sided setup, in our asset market setup, agents who exit may still trade—an agent who exits an ascending (descending) clock auction may trade as a seller (buyer). Thus, the usual clock auction interpretation needs to be adapted to this environment. We refer to our auction implementation as a *partial clock auction.*

In an ascending partial clock auction, we first increase the buyer clock price from zero until only $\frac{n-1}{2}$ active buyers remain. At this point, the set of agents with types that are less than or equal to $\hat{\theta}$, and their associated types, are observed. This allows us to initialize the state $\boldsymbol{\omega}_t = (\mathcal{B}_t, \mathcal{S}_t, p_t^B, p_t^S)$, with components as before, by setting $\mathcal{S}_0 \equiv \{i \in \mathcal{N} \mid \theta_i < \hat{\theta}\}$, $\mathcal{B}_0 \equiv \{i \in \mathcal{N} \mid \theta_i > \hat{\theta}\}$, and $p_0^B = p_0^S \equiv \hat{\theta}$. Then the clock auction defined previously can be used—as the buyer price increases, only the types of agents that exit from the buyer set are revealed. The privacy of all trading buyers is preserved. Analogously, one could start with a seller clock that descends from a price of one to identify the cutoff type and the set of agents with types greater than or equal to that type. Then the auction procedure preserves the privacy of all trading sellers.

While partial clock auctions endow agents with dominant strategies, they do not in general endow them with obviously dominant strategies. However, the truncated game that starts when the market clearing price is passed endows the remaining active agents with obviously dominant strategies from that point onwards because these remaining active agents are now essentially single-unit traders and the problem has become effectively one-sided.

The partial clock auction defined here is reminiscent of the incentive auction employed by the U.S. Federal Communications Commission in 2017 to purchase spectrum licenses from television broadcasters and then sell them to mobile wireless providers (see, e.g., Leyton-Brown et al., 2017; Milgrom, 2017). In that case, a reverse auction determined prices at which broadcasters were willing to relinquish their spectrum usage rights, and a forward auction determined the prices at which wireless carriers could purchase flexible use wireless licenses, with an intermediate process for the reorganization and reassignment of spectrum licenses. Apart from the complexity inherent in the incentive auction due to the interference constraints, conceptually, the key difference between the two formats is that here agents' trading positions are determined endogenously, whereas the incentive auction required bidders to register as buyers or sellers.

5 Extensions

In this section, we provide extensions. In Section 5.1, we generalize the market maker's objective to allow weight on both profit and social surplus. In Section 5.2, we extend the analysis beyond agents drawing their types from the same distribution to allow a priori heterogeneous groups of agents. Section 5.3 extends to general endowments and demands.

5.1 Ramsey objective

Market makers and designers may have objectives other than profit. In the tradition of Ramsey pricing (see e.g. Wilson, 1993), one natural candidate for an objective function is a weighted sum of expected profit and expected social surplus. In that case, for some $\alpha \in [0, 1]$, the objective is to maximize the expected value of the weighted sum

$$\alpha(\text{profit}) + (1-\alpha)(\text{social surplus}),$$

with the expectation being taken with respect to the distribution F. The prior-free mechanism that is asymptotically optimal relative to this objective is obtained by simply augmenting the virtual type estimators with α as follows:

$$\phi(\boldsymbol{\omega}) = p^B - \alpha\,(a+1)\,\sigma^B(\boldsymbol{\omega}) \text{ and } \gamma(\boldsymbol{\omega}) = p^S + \alpha\,(a+1)\,\sigma^S(\boldsymbol{\omega}). \tag{6}$$

All results above then extend to this setting.

5.2 Price discrimination among heterogeneous groups

We now explore the extent to which the analysis and mechanism extend to allow for ex ante heterogeneous groups of traders, at the cost, of course, of sacrificing constrained efficiency across groups because of the discrimination.[12] This problem is more complicated in an asset market than in standard one-sided or two-sided mechanism design problems, where it is a priori clear for which types the individual rationality constraint binds (the lowest possible type for a buyer and the highest possible type for a seller). Because the worst-off types are the same in each group on each side of the market, the designer knows ex ante where to start the elimination and estimation—at the lowest (highest) possible type for buyers (sellers). This is what permits clock implementation of asymptotically optimal prior-free mechanisms in one-sided and two-sided allocation problems with and without heterogeneous groups under the assumption that the environment is regular (Loertscher and Marx, 2020a).

In contrast, as is apparent from the above analysis, in an asset market problem, individual rationality constraints typically bind in the interior of the type space. This is, of course, what prevents clock implementation even with a homogeneous population. Nevertheless, what permits an asymptotically optimal prior-free asset market mechanism with a homogeneous population is that one knows that the agent whose type is equal to the median does not trade in the optimal mechanism. Because the median is a distribution-free concept, with a homogeneous population one can first identify the median and then start eliminating and estimating from there. With het-

[12] As will become clear, within each group the mechanism remains constrained efficient in the sense that only the highest value agents buy and only the lowest value agents sell.

erogeneous groups, as shown below, this procedure works as long as the heterogeneity among groups with respect to their relative sizes and distributions is not too accentuated, with the result that agents with types equal to the median do not trade in the optimal mechanism regardless of which group they belong to.

Formally, for the purposes of this section, we assume that there are m groups of agents. Let $\mathcal{M} \equiv \{1, \ldots, m\}$, and for each $j \in \mathcal{M}$, let n_j denote the number of agents in group j. We assume that each agent in group j draws its type independently from the distribution F_j, where, as before, F_j is continuously differentiable with positive density f_j on $[0, 1]$. Let Φ_j and Γ_j denote the corresponding virtual value and virtual cost functions, assumed to be increasing on $[0, 1]$. When analyzing the limits as the numbers of agents goes to infinity, we assume that those numbers increase in fixed proportion to one another.

Let $\overline{\mathbf{p}}^B = (\overline{p}_j^B)_{j \in \mathcal{M}}$ and $\overline{\mathbf{p}}^S = (\overline{p}_j^S)_{j \in \mathcal{M}}$ denote the Bayesian optimal prices in the limit economy when all groups or markets are integrated. These prices satisfy, for all $j \in \mathcal{M}$,

$$\Phi_j(\overline{p}_j^B) = \Gamma_j(\overline{p}_j^S) \equiv V \tag{7}$$

and the market clearing condition

$$\sum_{j \in \mathcal{M}} \mu_j (1 - F_j(\overline{p}_j^B)) = \sum_{j \in \mathcal{M}} \mu_j F_j(\overline{p}_j^S), \tag{8}$$

where $\mu_j \equiv \frac{n_j}{\sum_{i \in \mathcal{M}} n_i} > 0$ is the probability that a randomly selected trader draws its type from the distribution F_j.

Prior-free mechanism with heterogeneous groups

We now define a prior-free asset market mechanism for the setup with groups. As in the case without heterogeneous groups, the mechanism is direct, elicits reports from all agents, and identifies the cutoff type $\hat{\theta}$ as the median report. We initialize buyer and seller prices for each group $j \in \mathcal{M}$ as $p_0^{s_j} = p_0^{b_j} = \hat{\theta}$.[13] We denote the estimated virtual value for the marginal buyer in group j by ϕ_j and the estimated virtual cost for the marginal seller in group j by γ_j, where the estimation procedure is as above,

[13] As before, the median report is well defined if $n = \sum_j n_j$ is odd, but one can easily accommodate an even number of agents n with $n \geq 4$ by defining $\hat{\theta}_1 \equiv \theta_{(n/2+1)}$ and $\hat{\theta}_2 \equiv \theta_{(n/2)}$, where $\hat{\theta}_1 \leq \hat{\theta}_2$, and initializing the auction phase with $p_0^{s_j} \equiv \hat{\theta}_1$ and $p_0^{b_j} \equiv \hat{\theta}_2$.

except that the estimates for a given agent are based only on the reports of other agents from the same group. For each $t \in \{0, 1, \dots\}$, follow an iterative process similar to that defined above for the setup without groups, but with a state amended to include the sets of active buyers and sellers in each group and the buyer and seller prices in each group. In step 1 of the auction phase, if the minimum value of $\phi_j(\boldsymbol{\omega}_t)$ across all groups j with at least one remaining active buyer is greater than or equal to the maximum value of $\gamma_j(\boldsymbol{\omega}_t)$ across all groups j with at least one remaining active seller, then the auction phase ends. If not, then the buyer price for the group with the lowest virtual value is increased to be the minimum value among the remaining active buyers in that group, and that buyer is removed from the set of active buyers in that group, and the seller price for the group with the greatest virtual cost is decreased to be the maximum cost among remaining active sellers in that group, and that seller is removed from the set of active sellers in that group. Then t is incremented by 1 and the procedure continues with an updated state.

Asymptotical optimality with heterogeneous groups

Denote by p^W the market clearing price in the limit large economy with heterogeneous groups. That is, p^W is such that $\sum_{j \in \mathcal{M}} \mu_j (1 - F_j(p^W)) = \sum_{j \in \mathcal{M}} \mu_j F_j(p^W)$.

Proposition 6 *If for all $j \in \mathcal{M}$,*

$$\Phi_j(p^W) < V < \Gamma_j(p^W), \tag{9}$$

then the prior-free asset market mechanism for groups defined above is asymptotically optimal.

Condition (9) implies that agents with types equal or close to the market clearing price in the large economy do not trade in the Bayesian optimal mechanism independently of the group to which they belong. It prevents "excessive" heterogeneity across groups in the sense that, even though there is discrimination across groups, the discrimination effect is weak enough that an agent with type p^W neither trades in the agent's own group when that group is treated as a standalone market nor in the integrated market. One can therefore use the price p^W to initiate estimation and elimination without eliminating the "wrong" agents. As for the case of homogeneous agents, one can apply Loertscher and Marx (2020a, Theorem 1) to show that this mechanism is asymptotically optimal, which is why we do not include an independent proof of Proposition 6.

The Bayesian optimality condition in the large market limit in standalone market j is that the prices p_j^B and p_j^S satisfy both

$$\Phi_j(p_j^B) = \Gamma_j(p_j^S) \equiv V_j \tag{10}$$

and the market-clearing condition

$$1 - F_j(p_j^B) = F_j(p_j^S). \tag{11}$$

Conditions (10) and (11) holding for every standalone market j, we let

$$\underline{V} = \min_{j \in \mathcal{M}} V_j \quad \text{and} \quad \overline{V} = \max_{j \in \mathcal{M}} V_j.$$

Lemma 1 $V \in [\underline{V}, \overline{V}]$.

Proof. See Appendix B.

Lemma 1 implies that a sufficient condition for (9) to hold is what we refer to as *Virtual Warasian gap overlap*, which is stated formally as Assumption 1:

Assumption 1 *For all $j \in \mathcal{M}$,*

$$\Phi_j(p^W) < \underline{V} \quad \text{and} \quad \overline{V} < \Gamma_j(p^W).$$

As an illustration of Assumption 1, consider a model with two groups with n_1 agents in group 1 with types drawn from $F_1(\theta) = \theta$ and n_2 agents in group 2 with types drawn from $F_2(\theta) = \theta^2$. Thus, for group 1, we have $p_1^W = 1/2$ and $V_1 = 1/2$, and for group 2, we have $p_2^W = \sqrt{1/2} \cong 0.7071$ and $V_2 = \frac{1}{4}\sqrt{21 - 3\sqrt{17}} \cong 0.7345$, which means that $\underline{V} = 1/2$ and $\overline{V} \cong 0.7345$. Letting $k \equiv n_1/n_2$, in the limit as n_1 and n_2 go to infinity in fixed proportion, there are a large number of agents with types drawn from $H(x) \equiv (kF_1(x) + F_2(x))/(k+1) = x(k+x)/(k+1)$, with the Walrasian price in the large market given by $p^W = H^{-1}(1/2) = \frac{1}{2}(\sqrt{2 + 2k + k^2} - k)$. Then for all $k > 0$, we have $\Phi_2(p^W) < \Phi_1(p^W) < \underline{V}$ and $\overline{V} < \Gamma_2(p^W) < \Gamma_1(p^W)$ as illustrated in Figures 3. Thus, Assumption 1 is satisfied.[14]

[14]If we adjust the example so that $F_2(\theta) = \theta^3$, then Assumption 1 is satisfied if and only if

(a) Walrasian price in the large market limit

(b) Illustration of Assumption 1

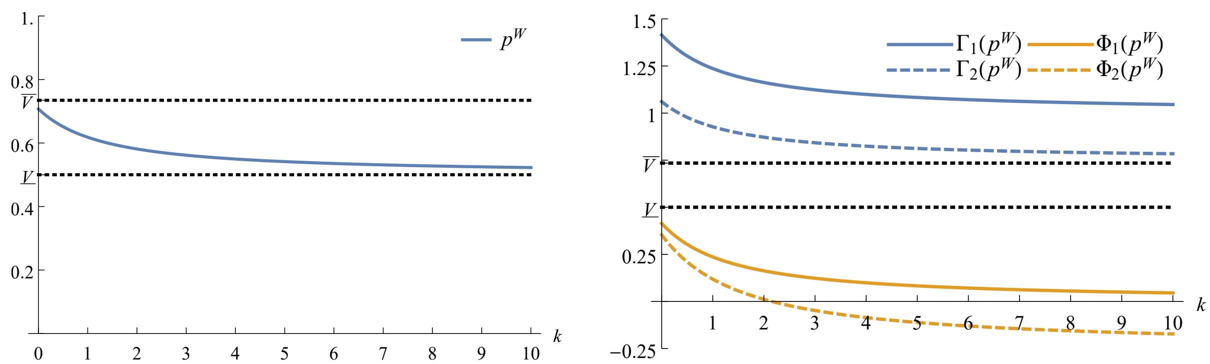

Figure 3: Walrasian price in the large and illustration of Assumption 1 for the case of two groups with $n_1 = kn_2$ for $k > 0$, $F_1(\theta) = \theta$, and $F_2(\theta) = \theta^2$.

Assumption 1, or even the condition that (9) holds for all $j \in \mathcal{M}$, could easily be relaxed by stipulating that they hold for some a priori *known* quantile q_j in every market j, that is, by requiring that for every $j \in \mathcal{M}$, there exists a $q_j \in [0, 1]$ such that $\Phi_j(F_j^{-1}(1 - q_j)) < \underline{V}$ and $\overline{V} < \Gamma_j(F_j^{-1}(q_j))$, respectively $\Phi_j(F_j^{-1}(1 - q_j)) < V < \Gamma_j(F_j^{-1}(q_j))$. To give an example, letting m_j be the median of distribution F_j, a sufficient condition is that for all $j \in \mathcal{M}$, $\Phi_j(m_j) < V < \Gamma_j(m_j)$.[15] Assumptions like these would allow all types that do not trade in the Bayesian optimal mechanism in the integrated market to vary across groups, that is, do not require that p^W is a type across all groups that does not trade in that mechanism. However, despite permitting additional flexibility and, thus, generality, the appeal of such alternative assumptions is admittedly limited because the idea that these quantiles are known ex ante may be difficult to reconcile with the idea of the mechanism being prior free.

5.3 General endowments and demands

We now show how the setup and analysis extends beyond the setting in which each agent has an endowment of one unit and a maximum demand for two.

$k \in (0.3125, 2.2865)$, i.e., the group sizes must be sufficiently similar for the assumption to hold.

[15]Noting that m_j is the market clearing price in standalone market j in the large limit, this condition could also be stated in terms of the standalone Walrasian prices.

Setup and Bayesian optimal outcome for the general case

In the generalized setup, each agent $i \in \mathcal{N}$ has endowment r_i and capacity k_i, where $0 \leq r_i < k_i$ for all $i \in \mathcal{N}$, with $r_i > 0$ for at least one $i \in \mathcal{N}$.

To define the Bayesian optimal outcome, we first introduce some notation. Define $\boldsymbol{\theta}^B$ to be the vector that replicates each agent i's type $k_i - r_i$ times:

$$\boldsymbol{\theta}^B \equiv (\overbrace{\theta_1, \ldots, \theta_1}^{k_1 - r_1 \text{ times}}, \ldots, \overbrace{\theta_n, \ldots, \theta_n}^{k_n - r_n \text{ times}}),$$

and define $\boldsymbol{\theta}^S$ to be the vector that replicates each agent i's type r_i times:

$$\boldsymbol{\theta}^S \equiv (\overbrace{\theta_1, \ldots, \theta_1}^{r_1 \text{ times}}, \ldots, \overbrace{\theta_n, \ldots, \theta_n}^{r_n \text{ times}}).$$

The counterpart to Lemma 2 is then:

Lemma 2 *In the Bayesian optimal outcome, q units are traded, where q is the largest index in*

$$\left\{0, 1, \ldots, \min\left\{\sum_{i \in \mathcal{N}} (k_i - r_i), \ \sum_{i \in \mathcal{N}} r_i\right\}\right\} \quad \text{such that} \quad \Phi(\theta^B_{(q)}) \geq \Gamma(\theta^S_{[q]}).$$

In the Bayesian optimal outcome, each agent i with type greater than $\theta^B_{(q)}$ purchases $k_i - r_i$ units, and so consumes its full capacity, and each agent i with type less than $\theta^S_{[q]}$ sells r_i and so consumes zero. The agent with type $\theta^B_{(q)}$ purchases $q - \sum_{i \text{ s.t. } \theta_i > \theta^B_{(q)}} (k_i - r_i)$ units and the agent with type $\theta^S_{[q]}$ sells $q - \sum_{i \text{ s.t. } \theta_i < \theta^S_{[q]}} r_i$ units.

To specify the threshold payments, first, we say that the seller side is the long side of the market if

$$\sum_{i \text{ s.t. } \theta_i \geq \theta^B_{(q)}} (k_i - r_i) \leq \sum_{i \text{ s.t. } \theta_i \leq \theta^S_{[q]}} r_i,$$

and otherwise we say that the buyer side is the long side of the market. If the seller side is the long side of the market, then all buyers pay $\theta^B_{(q+1)}$, which is less than $\theta^B_{(q)}$ when the seller side is the long side of the market. If the buyer side is the long side of the market, then all sellers are paid $\theta^S_{[q+1]}$, which is greater than $\theta^S_{[q]}$ when the buyer side is the long side. It then remains to describe the threshold payments for agents on the long side of the market, which are Vickrey prices (see, e.g., Krishna, 2002, Chapter 12). If the buyer side is the long side, then buyers' payments are the payments that would arise in a Vickrey auction to sell q units to the agents with types greater than or equal to $\theta^B_{(q)}$ with a reserve (minimum price) equal to $\max\{\theta_i \mid \theta_i < \theta^B_{(q)}\}$. If the seller side is the long side, then the payments to the sellers are the payments that would arise in a Vickrey auction to purchase q units from the agents with types less than or equal to $\theta^S_{[q]}$ with a reserve (maximum price) equal to $\min\{\theta_i \mid \theta_i > \theta^S_{[q]}\}$. For a review of the details of Vickrey pricing, see Appendix A or Loertscher and Marx (2020b).

Asymptotic optimality

In the generalized setup, asymptotics and asymptotic optimality are then captured, as in Gresik and Satterthwaite (1989), by looking at η-fold replicas of the economy characterized by $(k_i, r_i)_{i\in\mathcal{N}}$ and letting η go to infinity. More specifically, fixing the initial number of agents n, each replica adds an additional set of n agents characterized by $(k_i, r_i)_{i\in\mathcal{N}}$, with each agent independently drawing its type from F. Given $\eta \in \{1, 2, \dots\}$, an η-fold replica of this initial economy has ηn agents, where η have characteristics (k_1, r_1), η have (k_2, r_2), etc., up to η that have (k_n, r_n), each with an independently drawn type. We let $\mathcal{N}_\eta$ denote the set of agents in the η-th replica. In this setup, for all replicas, the designer's beliefs about the type space amount to beliefs over F, which is the same as for a single replica.

Let

$$\overline{r} = \sum_i \frac{r_i}{|\mathcal{N}|} \quad \text{and} \quad \overline{d} = \sum_i \frac{k_i - r_i}{|\mathcal{N}|}$$

be per-agent supply and net demand, respectively, in a given replica. Then, in the limit as η goes to infinity, the profit-maximizing mechanism is characterized by prices p^B and p^S such that agents with $\theta \in (p^S, p^B)$ do not trade, agents with $\theta_i \geq p^B$ buy $k_i - r_i$, and agents with $\theta_j \leq p^S$ sell r_j, with (p^S, p^B) satisfying

$$\Phi(p^B) = \Gamma(p^S)$$

and

$$\overline{d}(1 - F(p^B)) = \overline{r}G(p^S).$$

It follows that in the limit, the designer's per-agent expected profit is

$$\lim_{\eta \to \infty} \frac{1}{\eta n} \left(\sum_{i \in \mathcal{N}_\eta} \mathbb{E}\left[p^B (k_i - r_i) \mid \theta_i \geq p^B \right] - \sum_{i \in \mathcal{N}_\eta} \mathbb{E}\left[p^S r_i \mid \theta_i \leq p^S \right] \right).$$

Now that we have identified the Bayesian optimal outcome and characterized its asymptotic properties, we can define what it means for a mechanism to be asymptotically optimal. We say that an incentive compatible, individually rational mechanism defined for an η-fold replica is *asymptotically optimal* if the ratio of the designer's expected profit in the mechanism relative to its expected profit in the Bayesian optimal mechanism converges to 1 in probability as the number of replicas goes to infinity.

Incentive compatible asset market mechanism

The iterative steps in the auction phase of the asset market mechanism generalize as follows:

- Step 1: If $\phi(\boldsymbol{\omega}_t) \geq \gamma(\boldsymbol{\omega}_t)$ or $|\mathcal{B}| \leq 1$ or $|\mathcal{S}| \leq 1$, then the estimation phase ends. Otherwise, proceed to step 2.

- Step 2: Let

$$\boldsymbol{\omega}_{t+1} = \begin{cases} (\mathcal{B}_t \backslash \{j^B\},\ \mathcal{S}_t,\ \theta_{j^B},\ p_t^S) & \text{if } \sum_{i \in \mathcal{B}_t} (k_i - r_i) \geq \sum_{i \in \mathcal{S}_t} r_i, \\ (\mathcal{B}_t,\ \mathcal{S}_t \backslash \{j^S\},\ p_t^B,\ \theta_{j^S}) & \text{otherwise,} \end{cases}$$

  where j^B and j^S are defined by $\theta_{j^B} = \min\{\theta_i\}_{i \in \mathcal{B}_t}$ and $\theta_{j^S} = \max\{\theta_i\}_{i \in \mathcal{S}_t}$, and proceed to step 3.

- Step 3: Increment t by 1 and return to step 1.

When the estimation phase ends, with final state $\boldsymbol{\omega}^B = (\mathcal{B}, \mathcal{S}, p^B, p^S)$, then Q units trade, where

$$Q = \min \left\{ \sum_{i \in \mathcal{B}} (k_i - r_i),\ \sum_{i \in \mathcal{S}} r_i \right\}.$$

If $\sum_{i \in \mathcal{B}} (k_i - r_i) = \sum_{i \in \mathcal{S}} r_i$, then each buyer $i \in \mathcal{B}$ receives $k_i - r_i$ units and pays $p^B(k_i - r_i)$ and each seller $i \in S$ provides r_i units and is paid $p^S r_i$. If $\sum_{i \in \mathcal{B}} (k_i - r_i) < \sum_{i \in \mathcal{S}} r_i$, then each buyer $i \in \mathcal{B}$ receives $k_i - r_i$ units and pays $p^B(k_i - r_i)$ and the designer procures Q units from sellers in $\mathcal{S}$ using a Vickrey auction with reserve p^S. If $\sum_{i\in\mathcal{B}}(k_i - r_i) > \sum_{i\in\mathcal{S}} r_i$, then each seller $i \in \mathcal{S}$ provides r_i units and is paid $p^S r_i$ and the designer sells the Q units purchased from the sellers to the buyers in $\mathcal{B}$ using a Vickrey auction with reserve p^B.

With this generalization, Proposition 3 continues to hold. The proof follows as in the proof of Proposition 3, noting that the nature of clock and Ausubel auctions implies that in the auction phase every agent has a dominant strategy to bid sincerely.

Asymptotic optimality holds using the same virtual type estimators as defined above.

Heterogeneous groups

The setup of Section 5.2, where we allow heterogeneous groups of agents, extends to the agents with heterogeneous endowments and maximum demands if we adapt the market-clearing condition.

Letting $\overline{d}$ be the per-capita net demand and $\overline{r}$ be the per-capita supply, the market-clearing condition becomes

$$\overline{d} \sum_{j\in\mathcal{M}} \mu_j(1 - F_j(\overline{p}_j^B)) = \overline{r} \sum_{j\in\mathcal{M}} \mu_j F_j(\overline{p}_j^S). \tag{12}$$

With these adjustments, our prior results continue to hold. In particular, if for all groups j, $\Phi_j(p^W) < V < \Gamma_j(p^W)$ holds, where p^W is the market clearing price in the limit economy and V is the virtual type of the marginal traders in the Bayesian optimal mechanism in the limit, then the appropriately adjusted prior-free mechanism with groups is asymptotically optimal.

6 Conclusion

The advent of the gig economy—Uber, AirBnB, Nerdify, etc.—has rendered many markets that used to be two-sided into asset markets, with agents' trading positions being determined endogenously. Because these asset markets are typically organized by profit-seeking intermediaries who operate in new environments with demand and supply functions that are not known, they give renewed salience to the problem we studied here: developing prior-free asset market mechanisms that are asymptotically optimal.

Under the assumption that the underlying distributions are regular, we derive a prior-free, asymptotically optimal mechanism that allocates the quantity traded efficiently when all agents draw their types from the same distribution. With heterogeneous type distributions that differ across commonly known groups of agents, the mechanism generalizes, provided that the heterogeneity is not excessive. A major challenge in asset markets is identifying the types who do not trade in the profit-maximizing mechanism; these types are typically in the interior of the type space, making the problem different from one-sided or two-sided allocations problems (that is, auctions or double auctions), where it is a priori clear which buyer and seller types do not trade. Although privacy preservation via clock auctions is not possible in asset markets, our asymptotically optimal asset market mechanism can algorithmically be described as a clock auction and permits partial clock implementation. Further research on asset market problems with heterogeneous goods seems particularly relevant and promising.

A Appendix: Vickrey auction

This appendix draws heavily from Loertscher and Marx (2020b, Appendix A.1).

We first review a Vickrey auction that is used to *sell* q units with a reserve of p to a set of agents $\mathcal{A}$ with types that are all greater than or equal to p. The Vickrey auction allocates to each agent $i \in \mathcal{A}$ a quantity of $\overline{q}_i$ units, where

$$\overline{q}_i \equiv \min\{k_i - r_i, \;\; \max\{0, \; q - \sum_{\ell \in \mathcal{A} \text{ s.t. } \theta_\ell > \theta_i} (k_\ell - r_\ell)\}\}.$$

Thus, agent i is allocated units whenever there remain units available from the total of q after the net demands of agents in $\mathcal{A}$ with types greater than $\theta_i$have been satisfied, up to a maximum of $k_i - r_i$ units.

The amount paid by agent i is based on an individualized price vector consisting of p followed by the types of the other agents in $\mathcal{A}$, in increasing order by their types. Specifically, letting $\overline{\theta}^{-i}_{[\ell]}$ denote the type associated with the ℓ-th lowest type element of $\{\theta_j\}_{j \in \mathcal{A} \setminus \{i\}}$, we define:

$$\overline{\mathbf{p}}^i \equiv (p, \overline{\theta}^{-i}_{[1]}, \ldots, \overline{\theta}^{-i}_{[|\mathcal{A}|-1]}).$$

For example, in the setup of Figure 4, if $\mathcal{A} = \{1, 2, 3\}$ and $q = 3$, then the 3 units are allocated first to the higher-type agents so that $\overline{q}_1 = 1$, $\overline{q}_2 = 2$, and $\overline{q}_3 = 0$. Further, the individualized price vectors are simply the vectors consisting of p followed by the types of the other agents in $\mathcal{A}$, in increasing order: $\overline{\mathbf{p}}^1 = (p, \theta_3, \theta_2)$, $\overline{\mathbf{p}}^2 = (p, \theta_3, \theta_1)$, and $\overline{\mathbf{p}}^3 = (p, \theta_2, \theta_1)$.

	p	$<$	θ_3	$<$	θ_2	$<$	θ_1
$k_i - r_i$			1		2		1

Figure 4: Example of net demand by agents with types greater than p.

Agent i's total payment is determined by applying the prices in $\overline{\mathbf{p}}^i$ to traunches of agent i's units. Agent i pays p for the first $\overline{b}^i_0$ units that it purchases, pays $\overline{\theta}^{-i}_{[1]}$ for the next $\overline{b}^i_1$ units, etc. The traunches of units are defined so that each agent pays an amount equal to the externality that it exerts on the other agents. That is, letting $d^{-i}_{[\ell]}$ be the net demand (capacity minus endowment) of the agent associated with type $\overline{\theta}^{-i}_{[\ell]}$, we define $\overline{b}^i_0$ to be the maximum quantity (up to $\overline{q}_i$) out of q that can be allocated to agent i without affecting the allocation of the other agents in $\mathcal{A}$,

$$\overline{b}_0^i \equiv \max\left\{0,\ \min\left\{\overline{q}_i,\ q - \sum_{\ell=1}^{|\mathcal{A}|-1} d_{[\ell]}^{-i}\right\}\right\}, \tag{13}$$

and for $\ell \in \{1, \ldots, |\mathcal{A}| - 1\}$, we define $\overline{b}_\ell^i$ iteratively to be the additional quantity (up to $\overline{q}_i - \sum_{t=0}^{\ell-1} \overline{b}_t^i$) that can be allocated to agent i by imposing an externality on the ℓ-th lowest type agent in $\mathcal{A}$,

$$\overline{b}_\ell^i \equiv \max\left\{0,\ \min\left\{\overline{q}_i - \sum_{t=0}^{\ell-1} \overline{b}_t^i,\ d_{[\ell]}^{-i}\right\}\right\}. \tag{14}$$

Applying the prices in $\overline{\mathbf{p}}^i$ to the traunches for agent i, we get a total payment for agent i of $\overline{\mathbf{p}}^i \cdot \overline{\mathbf{b}}^i$.

In the example of Figure 4, assuming $q = 3$, we have $\mathbf{b}^1 = (0, 1, 0)$ and $\mathbf{b}^2 = (1, 1, 0)$. Thus, agent 1 pays θ_3 for its unit and agent 2 pays p for its first unit and θ_3 for its second unit. Because there are 3 units for sale and only two units demanded by agents other than agent 2, agent 2 is guaranteed to trade one unit and so pays the reserve for that unit. Agent 2's consumption of that unit has no impact on the trades of the others. However, agent 2's purchase of a second unit precludes agent 3 from trading, so agent 2 pays θ_3 for its second unit. Similarly, agent 1's purchase of a unit precludes agent 3 from trading, so agent 1 pays θ_3 for its unit.

Now turn to the Vickrey auction to purchase q units with reserve p from a set of agents $\mathcal{A}$ with types less than or equal to p. In this case, the quantity purchased from agent i is

$$\underline{q}_i \equiv \min\{r_i,\ \max\{0,\ q - \sum_{\ell \in \mathcal{A} \text{ s.t. } \theta_\ell < \theta_i} r_\ell\}\},$$

so that agent i sells units whenever there remain units demanded after the agents ranked below agent i have sold their endowments, up to a maximum of r_i units. Analogously to the buyer side case, let $\underline{\theta}_{(\ell)}^{-i}$ be the type associated with the ℓ-highest element of $\{\theta_j\}_{j \in \mathcal{A} \setminus \{i\}}$.

The individualized price vector for agent i is

$$\underline{\mathbf{p}}^i \equiv (p, \underline{\theta}_{(1)}^{-i}, \ldots, \underline{\theta}_{(|\mathcal{A}|-1)}^{-i}).$$

The payment made to agent i is determined by applying the prices in $\underline{\mathbf{p}}^i$ to traunches of units $\underline{\mathbf{b}}^i$ defined as in (13)–(14), but with $\overline{q}_i$ replaced by $\underline{q}_i$ and $d_{[\ell]}^{-1}$ replaced by the supply of the agent with type $\underline{\theta}_{(\ell)}^{-i}$. Then each agent $i \in \mathcal{A}$ is paid $\underline{\mathbf{p}}^i \cdot \underline{\mathbf{b}}^i$.

B Appendix: Proofs

Proof of Proposition 2. Let $x^* = p^S$ and note that because $y(x^*) = \Phi^{-1}(\Gamma(x^*))$, we have $y(x^*) = p^B$, giving us an ironing range of $[p^S, p^B]$. Define the allocation rule $\mathbf{s}^*$ so that when the empirical distribution of types is equal to F, which occurs with probability one in the limit, then

$$s_i^*(\boldsymbol{\theta}) = \begin{cases} -1 & \text{if } \theta_i < p^S, \\ 1 & \text{if } \theta_i > p^B, \\ 0 & \text{otherwise,} \end{cases}$$

which is feasible by (1). For other type realizations, let the allocation rule be (arbitrarily) such that no agent trades. Clearly, under this allocation rule and ironing region, the interim expected trade for agents with types in the ironing region is zero. With probability one, the empirical distribution is equal to F, so the Bayesian optimality of the allocation rule given the ironing region follows because for all $\theta_i > p^B$ and $\theta_j < p^S$, $\Phi(\theta_i) > \Gamma(\theta_j)$, and so $\mathbf{s}^*$ solves

$$\max_{\mathbf{s}} \mathbb{E}_{\boldsymbol{\theta}} \left[\lim_{n\to\infty} \sum_{i=1}^{n} \left(\Phi(\theta_i) s_i(\boldsymbol{\theta}) \cdot \mathbf{1}_{\theta_i > p^B} + \Gamma(\theta_i) s_i(\boldsymbol{\theta}) \cdot \mathbf{1}_{\theta_i < p^S} \right) \right],$$

subject to market clearing. ■

Proof of Proposition 4. Given Proposition 3, we can assume truthful reporting in the prior-free asset market mechanism.

Loertscher and Marx (2020a, Theorem 1) states that if F is continuously differentiable with positive density f and compact support and $b \in (0,1)$, as holds in our setup, then as $n \to \infty$,

$$\Pr\left[\sup_{\ell \in \{1,\dots,n-n^b\}} \left| \ell \frac{\theta_{(\ell)} - \theta_{(\ell+n^b)}}{n^b} - \frac{1 - F(\theta_{(\ell)})}{f(\theta_{(\ell)})} \right| > \varepsilon \right] \to 0 \tag{15}$$

and

$$\Pr\left[\sup_{\ell\in\{1,\ldots,n-n^b\}}\left|\ell\frac{\theta_{[\ell+n^b]}-\theta_{[\ell]}}{n^b}-\frac{F(\theta_{[\ell]})}{f(\theta_{[\ell]})}\right|>\varepsilon\right]\to 0. \tag{16}$$

At iteration $t\in\{0,1,\ldots\}$ of the prior-free asset market mechanism, we have state $\boldsymbol{\omega}_t$ with $a_t\equiv|\mathcal{B}_t|=|\mathcal{S}_t|=\frac{n-1}{2}-t$. Further, $\phi(\boldsymbol{\omega}_0)=\hat{\theta}-\left(\frac{n-1}{2}+1\right)\frac{1}{n+1}=\hat{\theta}-\frac{1}{2}$ and $\gamma(\boldsymbol{\omega}_0)=\hat{\theta}+\frac{1}{2}$, so $\phi(\boldsymbol{\omega}_0)<\gamma(\boldsymbol{\omega}_0)$, i.e., the stopping rule is never satisfied at $t=0$. Indeed, in the limit as n goes to infinity, with probability one, more that $n^b/2$ iterations are required prior to the stopping rule being triggered. At iteration t in the prior-free asset market mechanism, for $t>n^b/2$ we have

$$d(\boldsymbol{\omega}_t)=\min\left\{n^b,n-2a_t-1\right\}=\min\left\{n^b,2t\right\}=n^b$$

and spacing estimators

$$\sigma^B(\boldsymbol{\omega}_t)=\frac{\theta_{(a_t+1)}-\theta_{(a_t+1+n^b)}}{n^b}\equiv\sigma^B_{a_t+1}$$

and

$$\sigma^S(\boldsymbol{\omega}_t)=\frac{\theta_{[a_t+1+n^b]}-\theta_{[a_t+1]}}{n^b}\equiv\sigma^S_{a_t+1}.$$

Thus, $\phi(\boldsymbol{\omega}_t)=\theta_{(a_t+1)}-(a_t+1)\sigma^B_{a_t+1}$ and $\gamma(\boldsymbol{\omega}_t)=\theta_{[a_t+1]}-(a_t+1)\sigma^S_{a_t+1}$, and the algorithm ends with a_t trades (involving the a_t highest and lowest types) if

$$\theta_{(a_t+1)}-(a_t+1)\sigma^B_{a_t+1}\geq\theta_{[a_t+1]}-(a_t+1)\sigma^S_{a_t+1}.$$

In what follows, we first show that the allocation rule of the prior-free asset market mechanism based on these estimates converges in probability to the allocation rule of the Bayesian optimal mechanism. Then, second, we show that the ratio of the profit from using the prior-free asset market mechanism to the profit from the Bayesian optimal mechanism converges in probability to 1 by invoking the payoff-equivalence theorem and continuity to show that the probability that the per-buyer profits and the per-seller expenditures differ from those in the Bayesian optimum goes to zero as the number of agents goes to infinity.

As above, define p^S and p^B by $1-F(p^B)=F(p^S)$ and $\Phi(p^B)=\Gamma(p^S)$. By Proposition 2, in the limit as n goes to infinity, with probability one, the Bayesian optimal mechanism involves a purchase by every agent with a type greater than p^B

and a sale by every agent with a type less than p^S.

Given n, let $\tilde{\Phi}_n(j) \equiv \theta_{(j)} - j\sigma_j^B$ and $\tilde{\Gamma}_n(j) \equiv \theta_{[j]} + j\sigma_j^S$. Let i_n be the random variable such that when there are n agents, $i_n - 1$ agents buy and $i_n - 1$ agents sell in the prior-free asset market mechanism. By the definition of the prior-free asset market mechanism, i_n satisfies

$$\tilde{\Phi}_n(i_n) \geq \tilde{\Gamma}_n(i_n).$$

Using (15) and (16), for all $\varepsilon > 0$,

$$\Pr\left[\sup_{\ell \in \{1,\ldots,n-n^b\}} \left|\tilde{\Phi}_n(\ell) - \Phi(\theta_{(\ell)})\right| > \varepsilon\right] \to 0 \tag{17}$$

and

$$\Pr\left[\sup_{\ell \in \{1,\ldots,n-n^b\}} \left|\tilde{\Gamma}_n(\ell) - \Gamma(\theta_{[\ell]})\right| > \varepsilon\right] \to 0. \tag{18}$$

In other words, as $n \to \infty$, the probability that the thresholds for trade under the prior-free asset market mechanism and under the Bayesian optimal mechanism differ by more than ε, for any $\varepsilon > 0$, goes to zero. This allows us to prove the following lemma:

Lemma B.1 $\operatorname*{plim}_{n\to\infty} \theta_{(i_n)} = p^B$ *and* $\operatorname*{plim}_{n\to\infty} \theta_{[i_n]} = p^S$.

Proof of Lemma B.1. Let $\varepsilon > 0$ be given. By the definitions of p^B and p^S,

$$\Phi(p^B - \varepsilon) < \Phi(p^B) = \Gamma(F^{-1}(1 - F(p^B))) < \Gamma(F^{-1}(1 - F(p^B - \varepsilon))) \tag{19}$$

and

$$\Phi(p^B + \varepsilon) > \Phi(p^B) = \Gamma(F^{-1}(1 - F(p^B))) > \Gamma(F^{-1}(1 - F(p^B + \varepsilon))). \tag{20}$$

Suppose that

$$\operatorname*{plim}_{n\to\infty} \theta_{(i_n)} < p^B - \varepsilon, \tag{21}$$

which implies that $\operatorname*{plim}_{n\to\infty} \theta_{[i_n]} > F^{-1}(1 - F(p^B - \varepsilon))$. Using the continuity of Φ and Γ and Slutsky's Theorem, and the assumption that Φ and Γ are increasing, it follows that

$$\operatorname{plim} \Phi(\theta_{(i_n)}) = \Phi(\operatorname{plim} \theta_{(i_n)}) < \Phi(p^B - \varepsilon) \tag{22}$$

and

$$\Gamma(F^{-1}(1 - F(p^B - \varepsilon))) < \Gamma(\operatorname{plim} \theta_{[i_n]}) = \operatorname{plim} \Gamma(\theta_{[i_n]}). \tag{23}$$

By the definition of i_n, $\tilde{\Gamma}(i_n) \leq \tilde{\Phi}(i_n)$, so (17)–(18) imply that $\operatorname{plim} \Gamma(\theta_{[i_n]}) \leq \operatorname{plim} \Phi(\theta_{(i_n)})$. Combining this with (22) and (23), implies that $\Gamma(F^{-1}(1-F(p^B-\varepsilon))) < \Phi(p^B - \varepsilon)$, which contradicts (19), allowing us to conclude that (21) does not hold. Thus, we have $\operatorname{plim}_{n\to\infty} \theta_{(i_n)} \geq p^B - \varepsilon$. Analogously, using (20), we can conclude that $\operatorname{plim}_{n\to\infty} \theta_{(i_n)} \leq p^B + \varepsilon$, which, letting ε go to zero, completes the proof that $\operatorname{plim}_{n\to\infty} \theta_{(i_n)} = p^B$. Analogous arguments imply that $\operatorname{plim}_{n\to\infty} \theta_{[i_n]} = p^S$. □

Continuation of the proof of Proposition 4. The difference in the per-buyer expected payments between the prior-free asset market mechanism and the Bayesian optimal mechanism is, by the payoff-equivalence theorem,

$$\left| \int_{\theta_{(i_n)}}^{p^B} \Phi(x) f(x) dx \right|,$$

and similarly the difference in per-seller expected payment is

$$\left| \int_{\theta_{[i_n]}}^{p^S} \Gamma(x) f(x) dx \right|.$$

We are left to show that these differences converge in probability to zero. Fix $\Delta > 0$ and focus on the difference in per-buyer expected payments (the argument for per-seller payments is analogous). We need to show that

$$\lim_{n\to\infty} \Pr\left(\left| \int_{\theta_{(i_n)}}^{p^B} \Phi(v) f(v) dv \right| > \Delta \right) = 0.$$

To show this, note that for any $\varepsilon > 0$,

$$
\begin{aligned}
& \Pr\left(\left|\int_{\theta_{(i_n)}}^{p^B} \Phi(v)f(v)dv\right| > \Delta\right) && (24)\\
= \; & \Pr\left(\left|\int_{\theta_{(i_n)}}^{p^B} \Phi(v)f(v)dv\right| > \Delta \mid |\theta_{(i_n)} - p^B| \geq \varepsilon\right) \Pr\left(|\theta_{(i_n)} - p^B| \geq \varepsilon\right) && (25)\\
+ \; & \Pr\left(\left|\int_{\theta_{(i_n)}}^{p^B} \Phi(v)f(v)dv\right| > \Delta \mid |\theta_{(i_n)} - p^B| < \varepsilon\right) \Pr\left(|\theta_{(i_n)} - p^B| < \varepsilon\right) && (26)\\
\leq \; & \Pr\left(|\theta_{(i_n)} - p^B| \geq \varepsilon\right) + \Pr\left(\left|\int_{\theta_{(i_n)}}^{p^B} \Phi(v)f(v)dv\right| > \Delta \mid |\theta_{(i_n)} - p^B| < \varepsilon\right). && (27)
\end{aligned}
$$

The first term in the last line of the expression above is zero in the limit as n goes to infinity by Lemma B.1, and because $\Phi(v)$ and $f(v)$ are bounded, the second term is also zero in the limit as n goes to infinity because we can make ε arbitrarily small. Thus,

$$\lim_{n\to\infty} \Pr\left(\left|\int_{\theta_{(i_n)}}^{p^B} \Phi(v)f(v)dv\right| > \Delta\right) = 0$$

as required, and similarly for the seller side. ■

Proof of Proposition 5. In our setup, we have

$$\phi(\boldsymbol{\omega}_t) = p_t^b - (a_t + 1)\sigma^B(\boldsymbol{\omega}_t),$$

where

$$\sigma^B(\boldsymbol{\omega}_t) = \frac{\theta_{(a_t+1)} - \theta_{(a_t+1+d(\boldsymbol{\omega}_t))}}{d(\boldsymbol{\omega}_t)}$$

and

$$d(\boldsymbol{\omega}_t) = \min\left\{n^b, n - 2a_t - 1\right\}.$$

Note that $p_t^b = \theta_{(a_t+1)}$ and, if $d(\boldsymbol{\omega}_t) = n - 2a_t - 1$, then $p_t^s = \theta_{(a_t+1+d(\boldsymbol{\omega}_t))}$. In this case, $\sigma^B(\boldsymbol{\omega}_t) = \frac{p_t^b - p_t^s}{n-2a_t-1}$. Noting further that $a_t = \frac{n-1}{2} - t$, we have $\sigma^B(\boldsymbol{\omega}_t) = \frac{p_t^b - p_t^s}{2t}$. So, when $d(\boldsymbol{\omega}_t) = n - 2a_t - 1$, we have

$$\phi(\boldsymbol{\omega}_t) = p_t^b - \left(\frac{n+1-2t}{4t}\right)\left(p_t^b - p_t^s\right)$$

and, analogously,

$$\gamma(\boldsymbol{\omega}_t) = p_t^s + \left(\frac{n+1-2t}{4t}\right)\left(p_t^b - p_t^s\right).$$

We have $\phi(\boldsymbol{\omega}_t) \geq \gamma(\boldsymbol{\omega}_t)$ if and only if

$$p_t^b - \left(\frac{n+1-2t}{4t}\right)\left(p_t^b - p_t^s\right) \geq p_t^s + \left(\frac{n+1-2t}{4t}\right)\left(p_t^b - p_t^s\right),$$

which "first" occurs at the smallest t such that

$$1 - \frac{n+1-2t}{4t} \geq \frac{n+1-2t}{4t},$$

i.e., the smallest t such that $t \geq \frac{n+1}{4}$, i.e., $\hat{t} = \left\lceil \frac{n+1}{4} \right\rceil$. This means that the number of trades is $\frac{n-1}{2} - \left\lceil \frac{n+1}{4} \right\rceil$. If $2\left\lceil \frac{n+1}{4} \right\rceil \leq n^b$, then $d(\boldsymbol{\omega}_t) = n - 2a_t - 1$ for all $t \leq \left\lceil \frac{n+1}{4} \right\rceil$, so the stopping condition is satisfied before the number of spacings to be included in the average reaches n^b. ■

Proof of Lemma 1. Suppose to the contrary that the lemma is not true. That is, suppose, for example, that $V > \overline{V}$. By the monotonicity of the virtual type functions, this implies that for all $j \in \mathcal{M}$, $\overline{p}_j^B > p_j^B$ and $\overline{p}_j^S > p_j^S$, and hence $1 - F_j(\overline{p}_j^B) < 1 - F_j(p_j^B)$ and $F_j(\overline{p}_j^S) > F_j(p_j^S)$ for all $j \in \mathcal{M}$, which means that the market clearing condition (8) is violated. An analogous argument applies to the case $V < \underline{V}$, which proves the result. ■

References

Baliga, S. and R. Vohra (2003): "Market Research and Market Design," *Advances in Theoretical Economics*, 3, Article 5.

Baranov, O., C. Aperjis, L. M. Ausubel, and T. Morrill (2017): "Efficient Procurement Auctions with Increasing Returns," *American Economic Journal: Microeconomics*, 9, 1–27.

Chen, Y.-C. and J. Li (2018): "Revisiting the Foundations of Dominant-Strategy Mechanisms," *Journal of Economic Theory*, 178, 294–317.

Delacrétaz, D., S. Loertscher, and C. Mezzetti (2020): "When Walras Meets Vickrey," Working Paper, University of Melbourne.

Gresik, T. and M. Satterthwaite (1989): "The Rate at which a Simple Market Converges to Efficiency as the Number of Traders Increases: An Asymptotic Result for Optimal Trading Mechanisms," *Journal of Economic Theory*, 48, 304–332.

Johnson, T. R. (2019): "One-Sided Markets," Working Paper, University of Notre Dame.

Krishna, V. (2002): *Auction Theory*, Elsevier Science, Academic Press.

Lewis, T. R. and D. E. M. Sappington (1989): "Countervailing Incentives in Agency Problems." *Journal of Economic Theory*, 49, 294–313.

Leyton-Brown, K., P. Milgrom, and I. Segal (2017): "Economics and Computer Science of a Radio Spectrum Reallocation," *Proceedings of the National Academy of Sciences of the United States of America*, 114, 7207–7209.

Li, J. and P. Dworczak (2020): "Are Simple Mechanisms Optimal when Agents are Unsophisticated?" Working Paper, Northwestern University.

Li, S. (2017): "Obviously Strategy-Proof Mechanisms," *American Economic Review*, 107, 3257–3287.

Loertscher, S. and L. M. Marx (2020a): "Asymptotically Optimal Prior-Free Clock Auctions," *Journal of Economic Theory*, 187.

——— (2020b): "A Dominant-Strategy Asset Market Mechanism," *Games and Economic Behavior*, 1–15.

Loertscher, S. and C. Mezzetti (forth.): "A Dominant Strategy, Double Clock

Auction with Estimation-Based Tâtonnement," *Theoretical Economics*.

LOERTSCHER, S. AND C. WASSER (2019): "Optimal Structure and Dissolution of Partnerships," *Theoretical Economics*, 14, 1063–1114.

LU, H. AND J. ROBERT (2001): "Optimal Trading Mechanisms with Ex Ante Unidentified Traders," *Journal of Economic Theory*, 97, 50–80.

MILGROM, P. (2017): *Discovering Prices: Auction Design in Markets with Complex Constraints*, New York: Columbia University Press.

MILGROM, P. AND I. SEGAL (2020): "Clock Auctions and Radio Spectrum Reallocation," *Journal of Political Economy*, 128, 1–31.

SEGAL, I. (2003): "Optimal Pricing Mechanisms with Unknown Demand," *American Economic Review*, 93, 509–529.

SILVERMAN, B. W. (1986): *Density Estimation for Statistics and Data Analysis*, Chapman and Hall.

WILSON, R. (1987): "Game Theoretic Analysis of Trading Processes," in *Advances in Economic Theory*, ed. by T. Bewley, Cambridge, U.K.: Cambridge University Press, 33–70.

——— (1993): *Nonlinear Pricing*, Oxford University Press.